

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 08-08-2011		2. REPORT TYPE Proceedings		3. DATES COVERED (From - To) AUG 2011 - SEPT 2011	
4. TITLE AND SUBTITLE Hyperspectral Detection and Discrimination Using the ACE Algorithm				5a. CONTRACT NUMBER FA8720-05-C-0002	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) M. L. Pieper, D. Manolakis, R. Lockwood, T. Cooley, P. Armstrong, and J. Jacobson				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Kirkland AFB				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT One of the fundamental challenges for a hyperspectral imaging system is the detection and discrimination of subpixel objects in background clutter. The background surrounding the object, which acts as interference, provides the major obstacle to successful detection and discrimination. In many applications we look for a single signature and discrimination among different signatures is not required. However, there are important applications where we are interested for multiple signatures. In these cases, the use of spectral discrimination algorithms is both necessary and valuable. In this paper, we develop an approach to spectral discrimination based on the adaptive cosine estimation (ACE) algorithm. The basic idea is to jointly exploit the detection statistics from the various signatures and set a common threshold that ensures larger separation between signatures of interest and background. The operation of the proposed detection-discrimination approach is illustrated using real-world hyperspectral imaging data.					
15. SUBJECT TERMS Hyperspectral imaging, target detection, CFAR processing, matched filtering					
16. SECURITY CLASSIFICATION OF: U			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON Zach Sweet
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code) 781-981-5997

Hyperspectral Detection and Discrimination Using the ACE Algorithm

M. L. Pieper^a, D. Manolakis^a, R. Lockwood^a, T. Cooley^b, P. Armstrong^c, and J. Jacobson^d

^aMIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02115

^bSpace Vehicles Directorate Air Force Research Laboratory, 2251 Maxwell Avenue, Kirtland AFB, NM 87117

^cSpace Vehicles Directorate Air Force Research Laboratory, 29 Randolph Road, Hanscom AFB, MA 01731-3010

^dNational Air and Space Intelligence Center, Wright-Patterson AFB, OH 45433

ABSTRACT

One of the fundamental challenges for a hyperspectral imaging system is the detection and discrimination of subpixel objects in background clutter. The background surrounding the object, which acts as interference, provides the major obstacle to successful detection and discrimination. In many applications we look for a single signature and discrimination among different signatures is not required. However, there are important applications where we are interested for multiple signatures. In these cases, the use of spectral discrimination algorithms is both necessary and valuable. In this paper, we develop an approach to spectral discrimination based on the adaptive cosine estimation (ACE) algorithm. The basic idea is to jointly exploit the detection statistics from the various signatures and set a common threshold that ensures larger separation between signatures of interest and background. The operation of the proposed detection-discrimination approach is illustrated using real-world hyperspectral imaging data.

Keywords: Hyperspectral imaging, target detection, CFAR processing, matched filtering

1. INTRODUCTION

There are a vast number of applications where both military and civilian organizations want to remotely study a section of the earth. There may be objects in the scene that need to be identified, or maybe the background of the scene is of importance. Hyperspectral imaging is the detection and study of objects and materials from their spectral characteristics. A hyperspectral imaging sensor measures the electromagnetic spectrum of a spatial scene by dividing it into pixels like a digital camera. Each pixel consists of a spectrum which represents the materials present in the pixel. The spectra are created by measuring the electromagnetic spectrum at hundreds of narrow spectral wavelength bands. The result is a data cube with two spatial and one spectral dimension. The hyperspectral sensors that generated the data for this paper are passive sensors which measure the electromagnetic spectrum from 400 nm to 2500 nm. This wavelength range contains the visible up to the short wave infrared portions of the electromagnetic spectrum.

Detection algorithms are used to identify objects of interest in a scene, by comparing each pixel's spectrum to a known library spectrum for the object. There are a multitude of factors that determine how difficult detecting an object can be, but the main factors are spectral variability and background interference. The library spectra are only an approximation of how the object spectra will appear in a scene. Factors such as weathering, atmospheric propagation, and sensor noise effect the in-scene spectral signature. Other difficulties are how distinct the object's spectrum is relative to the background. If an object spectrum has a close resemblance to its surroundings, it will

Correspondence to M. L. Pieper E-mail: mpieper@ll.mit.edu, Telephone: 781-981-4085

This work is sponsored by the Air Force Research Laboratory under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

be more difficult to identify. Hyperspectral imagers have pixel sizes that can be several meters, and an object might only fill a fraction of the pixel. For subpixel objects the resulting spectrum is a linear combination of the object spectrum and the background spectrum weighted by their fill factors. As the size of the object decreases, the pixel appears more like the background, making it more difficult to identify the object.

In several applications there may be several objects of interest in a scene. In these situations, the signature library consists of the spectra for all the objects of interest. A detector is run for each object signature, to determine how similar each pixel is to each of the object signatures. Making the assumption that only one object can occupy a pixel at a time, a decision needs to be made on what object occupies the pixel. This is accomplished by assigning the pixel to the most similar object; this process is known as discrimination.

The purpose of this paper is to develop an approach to spectral discrimination based on the adaptive cosine estimation (ACE) detector. Existing methods treat each detector from a signature library independently, and set an independent Constant False Alarm Rate (CFAR) threshold for each detector. The basic idea of the proposed approach is to jointly exploit the detection statistics from the various signatures and set a common threshold that ensures that the CFAR threshold that is set is the actual CFAR of the result. The operation of the proposed detection-discrimination approach is illustrated using real hyperspectral imaging data.

The paper is organized as follows. Section 2 describes the linear background replacement model. Section 3 reviews the adaptive cosine estimator (ACE) and its statistical properties. Section 4 describes how ACE can be used as a discriminator. Section 5 describes two ACE detection and discrimination pipelines. The first pipeline sets a CFAR threshold for each object detector, and discriminates the results. The second pipeline discriminates every pixel, and sets a CFAR threshold for the discriminated results. Section 6 describes a data set that is used to validate the processing methods. This data set will be called Dataset A. Section 7 introduces the benefits of ACE discrimination thresholding using Dataset A. Section 8 shows how a library of signatures effects ACE discrimination thresholding. Section 9 lists discrimination and false alarm results for each of the processing methods on Dataset A. Section 10 shows the benefits of ACE discrimination thresholding using two other datasets collected with a different sensor.

2. SIGNAL MODEL

Hyperspectral scenes consist of background materials and objects of interest. The background usually consists of several materials. These materials can include vegetation such as grass and tree tops, roads, dirt, and building tops. Each of these materials is labeled as a background class, and is assumed to be normally distributed. In general, when we try to detect an object in an arbitrary scene, the background composition of the scene is usually unknown. There are several ways to cluster a scene, and split the background into its different background classes. However, the clustering methods are computationally intensive and take more time than the detection itself. For this reason, a scene is usually split into two classes, background and object classes. The background is still assumed to be normally distributed, even though it is typically multimodal.

With only one background class, all the spectra consist of a linear mixture of the a random background-plus-noise spectrum \mathbf{v} and a target of interest spectrum \mathbf{t} . When the target is opaque, it replaces the background within a pixel. The fraction of background the object replaces a , corresponds to the area of the pixel that it occupies. This linear mixture is described by¹

$$\mathbf{x} = a\mathbf{t} + (1 - a)\mathbf{v}, \quad 0 \leq a \leq 1. \quad (1)$$

When detecting objects in a scene, a decision has to be made for every pixel. This decision involves deciding between one of two hypotheses. The first hypothesis is that the pixel contains only background, and the second hypothesis is that the pixel contains a mixture of background and an object of interest.

$$\begin{aligned} H_0 : & \quad \mathbf{x} = \mathbf{v} \\ H_1 : & \quad \mathbf{x} = a\mathbf{t} + (1 - a)\mathbf{v}, \quad 0 \leq a \leq 1. \end{aligned} \quad (2)$$

A detection algorithm uses a detection statistic and a CFAR threshold to determine the presence or absence of an object.

3. ADAPTIVE COSINE ESTIMATOR (ACE)

The coherent ACE detector is defined by²

$$y = \frac{(s - \mu_b)^T \hat{\Sigma}^{-1} (x - \mu_b)}{\sqrt{(s - \mu_b)^T \hat{\Sigma} (s - \mu_b)} \sqrt{(x - \mu_b)^T \hat{\Sigma}^{-1} (x - \mu_b)}}. \quad (3)$$

The ACE essentially operates in a zero-mean whitened space. In this whitened space the ACE is the cosine of the angle between the demeaned pixel spectrum and the demeaned object spectrum. In this paper, we will be dealing with solid targets. It is important that the pixel spectrum is pointing in the same direction as the target spectrum. The incoherent ACE which is the square cosine of the angle between the target and signature would cancel out the direction sign. For this reason, the coherent ACE is used. Any pixel that is pointing in the opposite direction of the object spectrum will be negative, and will be easily removed. This is especially helpful when a scene contains objects whose spectra have a large angle with the object of interest.

The most significant characteristic of the coherent ACE is that the distribution of the background is a central t-distribution. The standard deviation of this t-distribution is directly related to the square root of the number of degrees of freedom, or the number of usable bands. Using this information, a CFAR threshold can be set based on the background distribution.³

One implication of the standard deviation being related to the number of bands, is that the amount of separation between possible objects of interest and the background is directly dependent on the number of bands. The ACE detector can have values between -1 and $+1$. Full object pixels will have a value around 1 , but the ACE value will decrease with the fill fraction. At some point, the ACE value for a subpixel object will be small enough where it is swallowed by the background. The more quality bands that are used, the smaller the standard deviation, and the narrower the background distribution. With a narrower background distribution, one will be more likely to detect small objects without significant false alarms.

The denominator of the ACE equation is the Mahalanobis Distance

$$m = \sqrt{(x - \mu_b)^T \hat{\Sigma}^{-1} (x - \mu_b)} \quad (4)$$

The Mahalanobis distance is a distance measure similar to the Euclidean distance, however, it takes into account the correlations of the data set and is scale-invariant. The Mahalanobis distance is used as an anomaly detector. Anomalies include pixels containing non-background materials. The Mahalanobis distance is compared to a threshold and used to remove non-background materials from the scene.

4. ACE DISCRIMINATION

The ACE statistic is a measurement of the likelihood that a pixel contains an object of interest. If we look for multiple objects of interest, an ACE can be run for each of the object signatures. The results for each of the detectors, will determine the likelihood that each of the pixels contains the corresponding object.

Figure 5 is an illustration of how ACE can be used to discriminate between object signatures. In the illustration all of the vectors are in a whitened zero-mean space. In this space, ACE is the angle between the pixel-under-study (PUT) and the object of interest. It can be seen that there is a smaller angle between \mathbf{s}_1 and the PUT. With a smaller angle, the PUT would be classified as \mathbf{s}_1 . The ACE detector is a superior discriminator than the matched filter. The matched filter is the projection of the pixel onto the object signature. The matched filter response is shown for each of the object signatures. It can be observed that matched filter responses for each of the signatures are fairly similar, where the ACE responses were not.

When the ACE detector is used to discriminate between object signatures in this way, it is assumed that only one object of interest can be occupying a pixel at a time. The vector that connects the background to the object signature, is the vector where the background and object mixtures lie. The matched filter is a representation of the objects fill fraction within a pixel. The ACE is a measurement of how well a given pixel fits the mixing model. If one wanted to detect pixels that contained several objects of interest, a subspace detector would have




Figure 1. Illustration of the ACE response being used to discriminate which object is present within a pixel.

to be used. This subspace would consist of the demeaned signatures of the objects you are looking for. There are several problems with using a subspace for detection purposes. The most significant problem is as you increase the size of the subspace, you will retrieve more pixels and obtain more false alarms. The other problem is even with a relatively small number of signatures, the number of combinations of signatures will increase quickly. The likelihood that there are combinations of signatures in a pixel, is directly dependent on the relative size of the objects compared to the pixel size of the sensor. A larger pixel size compared to the objects of interest, means that there is more of a likelihood that multiple objects will occupy a single pixel. On the other hand, if the pixels have one quarter the square area of the square objects of interest, it is guaranteed that the objects will entirely occupy a pixel.

5. ACE DETECTION AND DISCRIMINATION PROCESSING

Detecting and discriminating between several object signatures is a multi stage process. Depending on whether you are working on a radiance cube or a reflectance cube, it is important that your signature library is in the same units as the data cube. In this paper, all the data was converted to reflectance using the empirical line method (ELM).⁴ To run a detection algorithm on the data cube, the next step is to estimate the background statistics of the scene. To prevent outliers from corrupting the covariance matrix, all the outliers are removed. This process is known as target free background estimation (TFBE). TFBE is a two stage anomaly detection process that uses the Mahalanobis distance detector. The first stage runs the Mahalanobis distance detector using the mean and covariance of the entire scene. Pixels that are a specific number of standard deviations from the mean are removed, and the process is repeated. After the second run through, all the outliers have been removed from the scene and a target free background mask has been created. The problem with using the Mahalanobis distance as an anomaly detector is it does not remove subpixel objects that consist mostly of background. By using the Mahalanobis distance as a metric, subpixel objects are close to the background because they consist mostly of background. However, the subpixel objects are in the same direction as the vector between the background and object signature, and they can have adverse effects on the background covariance. Assuming that there are a small number of subpixel objects in a scene, their effect on the background covariance should be minimal. Large objects will have mostly mixed pixels with small fill fractions along their edges. If the fill fractions of these pixels are small enough, they will not be removed from the scene. To remove mixed pixels along the edges of large objects, the original target free background mask is dilated. As long as there are enough background pixels to calculate a covariance that is non-singular and can be inverted, it is beneficial to throw away any pixel that will corrupt the covariance.

Using the calculated background statistics, an ACE detector is run for every signature in the object library. The detectors for each signature show the likelihood that a pixel contains the object of interest. Using all this

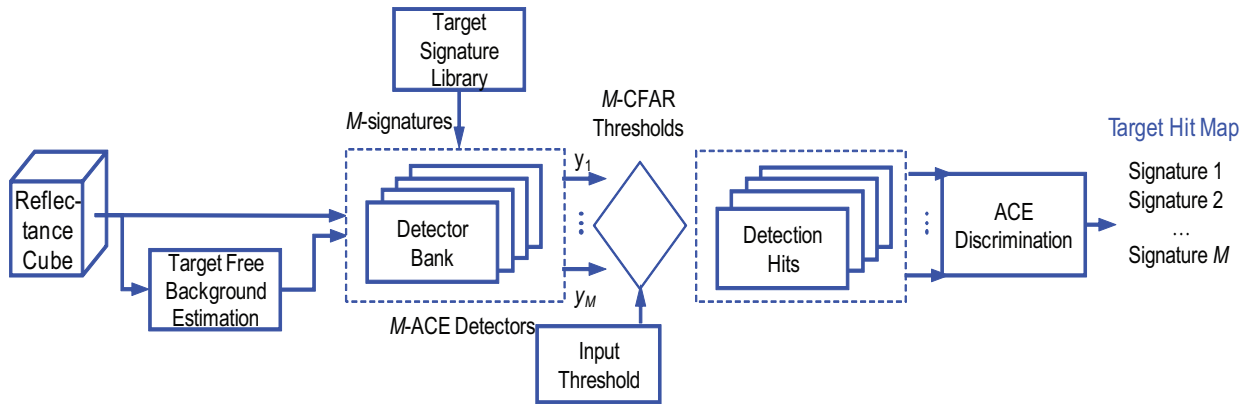


Figure 2. Illustration of the ACE response being used to discriminate which object is present within a pixel.

information it can be determined whether or not an object is present in a pixel, and then the type of object.

The next two subsections describe the detection and discrimination pipelines. Each pipeline begins with the target free background estimation and detection process that were previously described. The pipelines differ by the order of the discrimination and CFAR thresholding stages.

5.1 Signature Dependent Thresholding

The pipeline that uses signature dependent thresholding is shown in Figure 2. As stated above the first two stages of the pipeline are: target free background estimation, and running an ACE detector for every signature in the library.

The next step involves setting a CFAR threshold for each detector to split up the background pixels from the possible object pixels of interest. The benefit of signature dependent thresholding is the background distribution of the ACE detector is well known. As it was mentioned in section 3, the background distribution of the ACE detector is a central t-distribution. In the high dimensional space of the the hyperspectral data, the t-distribution becomes a Gaussian. Since, the standard deviation of the distribution is known to be one over the square root of the number of bands, one can set a CFAR threshold by using the CDF of a Gaussian with the known standard deviation. Using the CFAR thresholds, a set of detection hits is retrieved from each detector output.

Depending on the difference between the signatures in the object library, and the signatures of the objects in the scene, pixels may pass the thresholds of several of the detectors. Assuming that a pixel can only contain background and one object at a time, pixels that have hits from several ACE detectors are discriminated using the ACE responses to determine which object occupies the pixel. Discrimination can be thought of as a process of false alarm removal. If the object in the scene is one of the objects in the library, discrimination identifies the object correctly, and removes the pixel as a false alarm from the other detectors. If the object is not present in the library, it is identified by the closest signature and a false alarm will result. Removing such false alarms is beyond the scope of this paper. After discriminating, the result is a list of hits, and which library signature each hit corresponds to.

5.2 ACE Discrimination Thresholding

The ACE Discrimination Thresholding pipeline is nearly identical to the signature dependent thresholding pipeline. The difference between the two pipelines is that the order of the CFAR thresholding and ACE discrimination is switched. The new pipeline is shown in Figure 3. Since the two pipelines are identical up to the detection stage, the ACE responses for both of the pipelines will be identical. When the ACE responses are discriminated before thresholding, every pixel in the scene is being discriminated, this merges the background distributions for all the detectors by taking the maximum ACE response for each pixel. The merged background detections are no longer a central t-distribution. A discrimination log is also created which lists which object had the highest ACE response for each pixel. When a single CFAR threshold is set using the merged background

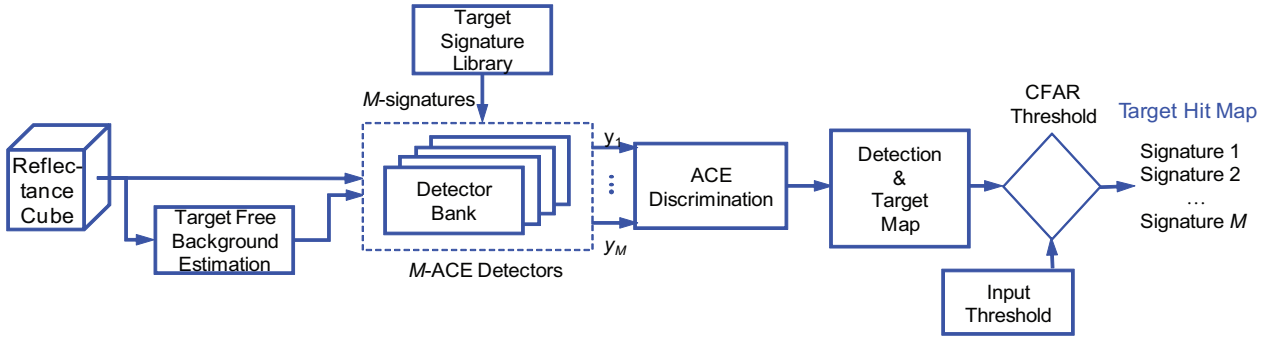


Figure 3. Illustration of the ACE response being used to discriminate which object is present within a pixel.

detection, the resulting threshold is an improvement upon the signature dependent thresholds. This threshold takes into account the detections from all of the detectors. After setting the threshold and obtaining the pixels containing objects of interest, the detection log is then used to determine which object occupies each of the pixels.

Since the merged background distribution is not a central t-distribution, and the distribution is slightly skewed to the right, the properties of a Gaussian distribution with a known standard deviation cannot be used to set a CFAR threshold. In order to set a CFAR threshold for the merged background distribution, a probability of exceedence plot is created for the the background pixels in the scene.

A probability of exceedence plot is a tool used to determine the value of the decision threshold. The probability of exceedence $P_y(\eta)$ is defined as the probability that the value of y , with probability density function $f_y(y)$, exceeds a threshold η . That is, the detection statistic

$$P_y(\eta) = \int_{\eta}^{\infty} f_y(y) dy = 1 - \int_{-\infty}^{\eta} f_y(y) dy = 1 - F_y(\eta), \quad (5)$$

where $F_y(\eta)$ is the cumulative distribution function.⁵ Any pixels with detection statistics greater than the threshold η are classified as object hits.

The TFBE mask is used to retrieve the background pixels; however this mask may contain subpixel objects of interest. To remove these pixels, an initial threshold is set for the ACE discriminated responses using the same method used for the TFBE. Any pixel in the TFBE that does not pass the initial threshold is thrown away as a subpixel object. The ACE discriminated responses for the background pixels are used to create a probability of exceedence plot. A CFAR is chosen, and the plot is used to determine the corresponding threshold.

6. DESCRIPTION OF DATA

The above pipelines were used to process Dataset A. Dataset A is a forest scene, with several objects of interest placed in a large field. The data was taken at a nadir viewing angle, at an altitude of 2000 meters. At this altitude, the approximate pixel size was 0.7m x 0.7m. The data was retrieved using the COMPASS sensor, which has a sample width of 254 pixels, and 255 spectral bands across the whole spectral range from 0.414 μm to 2.4 μm . Only 159 of the spectral bands were used after removing the atmospheric absorption and low-SNR bands.

The spectral library used to process the data is shown in Figure 4. The shaded out regions represent bands that have been removed. In total there are 8 signatures in the library. Looking at the signatures, it can be seen that signatures 2, 4, and 8 are very similar to each other. This can cause problems with object discrimination. Of the 8 object signatures, only signatures 2 through 5 have objects that are present in the scene. In total, there are 34 objects with ground truth that were used to validate the pipelines. Of the 34 objects, 20 are signature 2 and 10 are signature 3. The objects for both signatures 2 and 3 were all square with widths that varied between 1.7 m and 0.27 m. With these dimensions there are both full pixel and subpixel objects present in the scene. There is one object with signature 4, and three objects with signature 5. The objects with signature 4 and 5 are large objects that fully occupy several pixels.

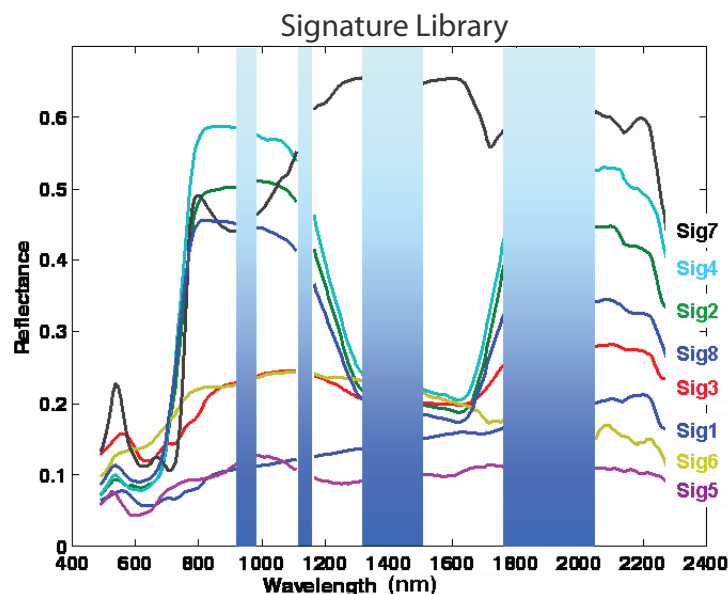


Figure 4. Object signatures in library.

7. ACE DISCRIMINATION THRESHOLDING BENEFITS

There is a significant difference between setting a CFAR threshold for every individual detector, and setting a CFAR threshold for the merged detector results. The last stage of the of the ACE Discrimination Thresholding pipeline is the setting of the CFAR threshold. By setting the threshold last, the distribution of the data is based on the discrimination. The results for both of the pipelines are based on discrimination. By using the distribution of the final result, one can accurately set a CFAR rate for the final result. This is described in Figure 5.

Figure 5 shows Probability of Exceedence plots of the background pixels for the individual ACE detectors in blue, and the ACE discriminated detector in red. It can be seen from the plots that the ACE discriminated detector always has a higher threshold at the same false alarm rate of the individual detectors. The green and magenta lines show CFAR threshold of 10^{-3} for the individual and ACE discriminated detectors respectively. By studying the points at which the lines pass the other detectors, one can observe what false alarm rate would have to be set to obtain the results of the other. To get the same results as the ACE discriminated detector, with the individual detectors, one would have to set a CFAR threshold for each detector at approximately 10^{-4} . The individual detectors CFAR of 10^{-3} is equivalent to approximately a 10^{-2} CFAR for the ACE discriminated detector. This means that the individual CFAR thresholds will pass an order of magnitude more background false alarms.

When a CFAR threshold is set for an individual detector, it is based on the angle between the object spectrum and the background pixels. By setting a CFAR threshold you are statistically passing a fraction of the outlying background pixels that are closest to the object spectrum. The outlying background pixels change for each object spectrum, due to the differences between the object spectra. By setting a CFAR threshold for each individual detector, you accumulate outlying background pixels which typically consist mostly of false alarms. This means that as you look for more objects, and increase the number of the detectors, you will accumulate more outlying background false alarms. When a CFAR threshold is set for the ACE Discrimination detector, all the outlying background pixels for each detector are conglomerated at the end of the merged background distribution. Using the merged background distribution allows one to set a single CFAR threshold for the entire detector set, and provide a result that will not accumulate false alarms based on the number of signatures in the library.

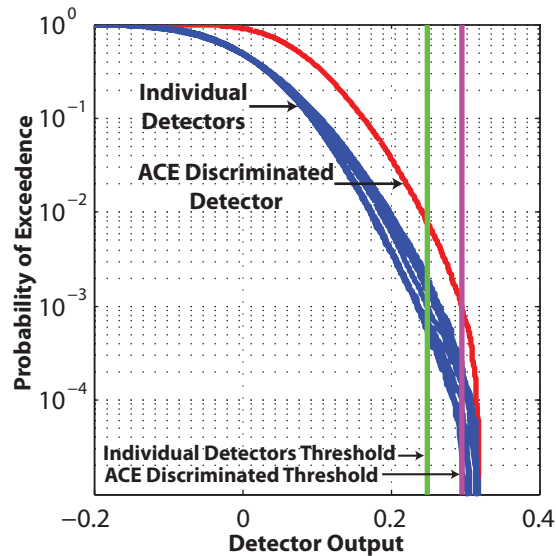


Figure 5. Probability of Exceedence of background for individual detectors, and merged detectors.

8. LIBRARY EFFECTS ON ACE DISCRIMINATION BACKGROUND

The background results for each of the detectors has a t-distribution whose standard deviation was dependent on the number of bands used. Even though each of the detectors has the same background distribution, the results are different for each of the detectors. The angle between a pixel and each of the object spectra differs, and this difference determines where the pixel lands in the detector's background distribution. By taking the maximum of all the detection results, a new background distribution is created. This distribution is how close each pixel is to the closest object of interest. The new background distribution depends on the objects of interest themselves. The greater the angles between the objects of interest the more the background distribution will differ from the original.

Figure 6 shows the angles between the whitened signatures in the Dataset A library. Signatures 2, 4, and 8 are all similar with angles less than 10 degrees between them. Signatures 6 and 7 share an angle of 60 degrees between them, and are nearly orthogonal to the rest of the signatures. Seeing mostly high angles in the picture means that there is a wide variation between the signatures, and that where a pixel lies in the central t-distribution will vary greatly between the signatures.

In Figure 5 it was shown that the ACE Discriminated background had a positively shifted distribution. The amount of the shift was dependent on the angles shared between the signatures shown in Figure 6. The angle between two signatures determines the correlation between their ACE responses. The larger the angle the lower the correlation. When two signatures share a large angle between them, pixels that are far from one signature tend to be closer to another. When the ACE responses for both these signatures are discriminated, you increase the mean response of the distribution by replacing the further responses with much closer ones. The lower the correlation is between the ACE responses, the larger the shift will be evident in the discriminated background distribution. Signatures with a high correlation will result in very little shift.

Figure 7 shows figures that describe how the relationship between two library signatures effects the ACE discriminated background. On the left are density scatter plots between the ACE responses of two sets of signatures for the background pixels in Dataset A. On the right are the ACE discriminated background histograms for the same two sets of signatures. The top pair of plots are for signatures 2 and 4. As it is shown in Figure 6 these two signatures share an angle of less than 10 degrees between them. This means that the two signatures have a high correlation coefficient, meaning there is a high linear dependence between the two ACE responses. This is shown in the density scatter plot, the points are packed against the black line. The black line in the scatter plot is a discrimination line, what signature a pixel is classified as is determined by which side of the




Figure 6. Shows the Mahalanobis angles between object Signatures for Dataset A.

discrimination line the pixel lies. Since the two signatures are so similar, not much information is acquired by background discrimination. For this reason, the ACE discriminated background histogram remains a central Gaussian distribution. The top pair of plots are for signatures 2 and 6. These two signatures are nearly orthogonal to each other with a correlation coefficient close to zero. This is shown in the density scatter plot, the points form a circle which is split by the discrimination line. The uncorrelation between the signatures causes the shift shown in the ACE discriminated background histogram. It should be noted that even though the mean of the distribution shifted, the upper bound of the distribution remained constant.

9. ACE DISCRIMINATION RESULTS

The ACE Discrimination Pipeline was used to process Dataset A. Figure 8 shows the detection and discrimination results. Each element of the table lists the detection vs the opportunities. The opportunities are the total number of pixels an object occupied in the scene. The detections describe how many of an object's pixels in the scene were classified as each object. The detections along the diagonal signify correct identifications. It can be observed that all the objects were identified correctly. There were some misclassification between signatures 2, 4, and 8, but these signatures are very similar to each other because they are the same paint that has weathered for varying amounts of time. The last row lists the number of false alarms for each of the signatures. In total, there were 192 false alarms.

The smallest signature 2 object had the lowest ACE response at 0.6883. This is far above the $1e-3$ CFAR threshold of 0.3195. If the original pipeline was used, the $1e-3$ CFAR thresholds would be between 0.2420 and 0.2783, with an average of 0.2586. Both of the pipelines are based on identical ACE detectors, and the only difference is when the CFAR threshold is set. The discrimination results are identical for the pixels that pass the thresholds for both pipelines. Since the lowest ACE response for an object of interest is above the thresholds for both pipelines, the discrimination results using the individual detector pipeline are identical to the results shown in Figure 8. However, the number of false alarms for the individual detector pipeline is 634 because the thresholds are actually at approximately a $1e-2$ CFAR.

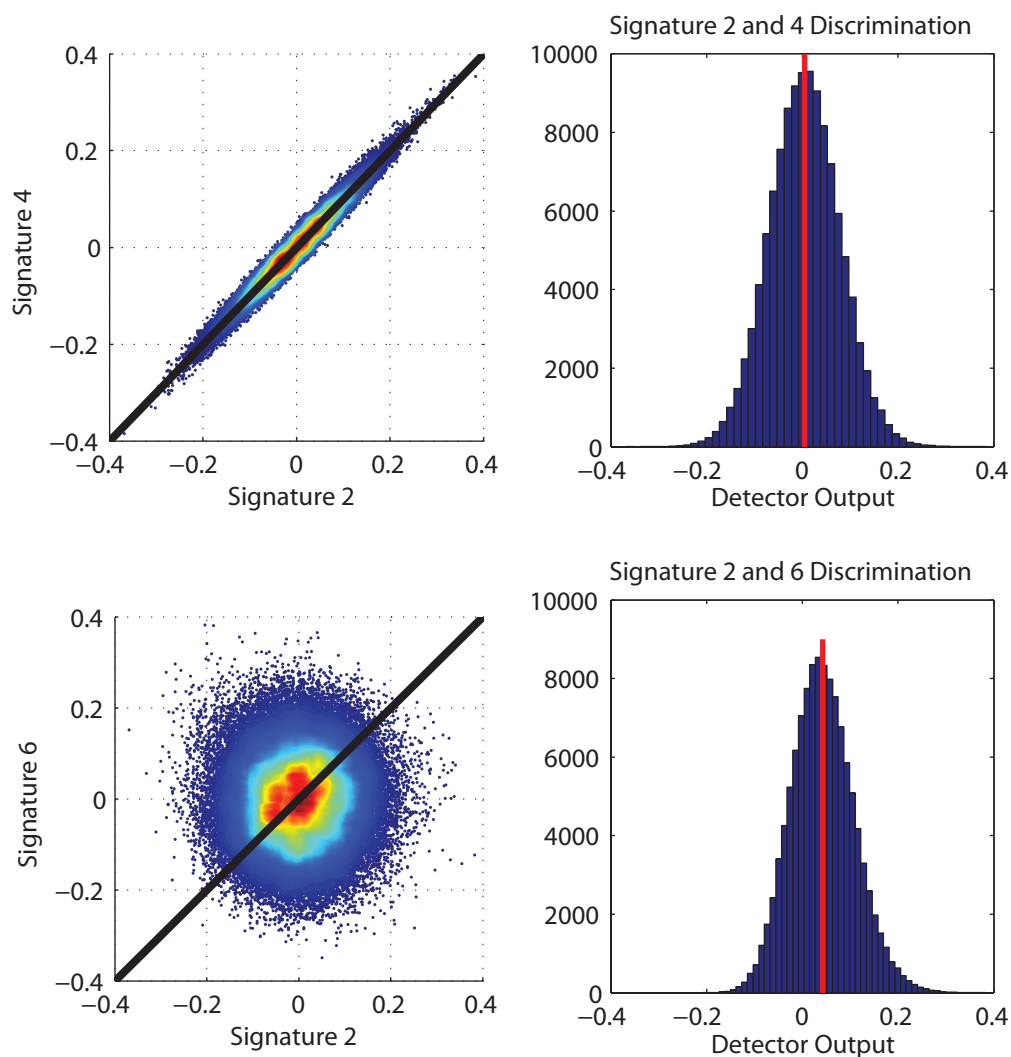


Figure 7. Left: ACE Response Density Scatter Plots. Right: ACE Discrimination Background Histograms.

Targets Present	Signatures of Interest (Detections/Opportunities)							
	Sig1	Sig2	Sig3	Sig4	Sig5	Sig6	Sig7	Sig8
Sig1	0/0							
Sig2		152/152						
Sig3			55/55					
Sig4		24/63		34/63				5/63
Sig5					164/164			
Sig6						0/0		
Sig7							0/0	
Sig8								0/0
FA	39/192	1/192	14/192	13/192	4/192	42/192	77/192	2/192

Figure 8. Discrimination Results for the merged detectors.

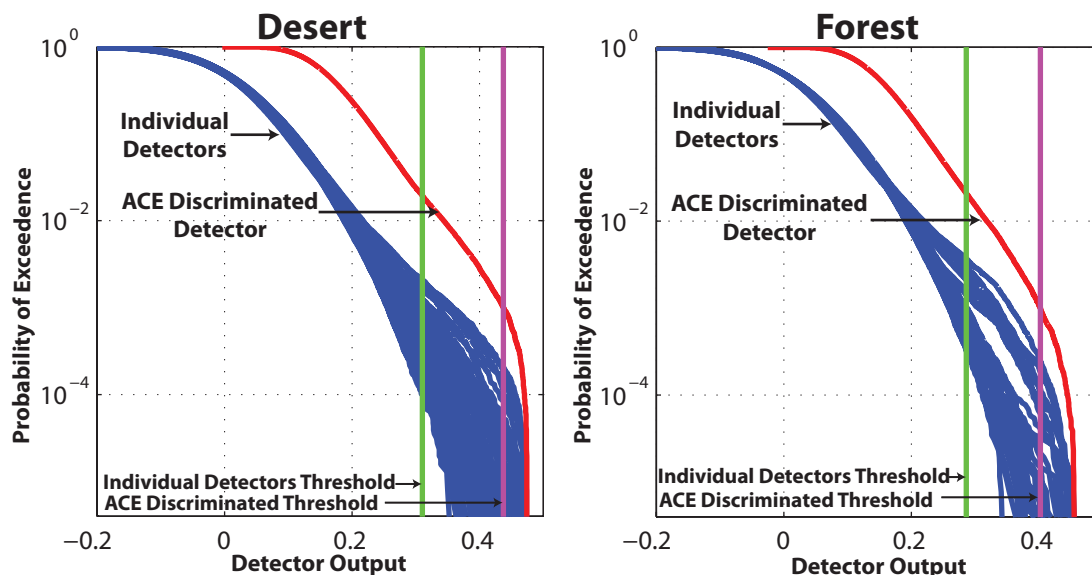


Figure 9. Probability of Exceedence of background for individual detectors, and merged detectors. Left: Dataset B Right: Dataset C

10. HYDICE RESULTS

The Signature Dependent and ACE Discrimination Thresholding pipelines were tested on two more scenes taken with the HYDICE sensor. Dataset B is a desert scene with dirt roads, and Dataset C is a forest scene with paved roads. These tests were conducted to show that the pipelines provide similar results with data taken with a different sensor, and consisting of varying backgrounds. Each scene was collected at an altitude of approximately 10000 ft. The sensor has a sample width of 306 quality pixels, and 210 spectral bands across the whole spectral range from $0.4 \mu\text{m}$ to $2.5 \mu\text{m}$. Only 145 of the spectral bands were used after removing the absorption and low-SNR bands. For each scene, 960 lines were collected. The desert scene had 65 signatures in its library, and the forest scene has 24 signatures. These signatures consisted of tan and green fabrics and paints, and several types of vehicles.

The two pipelines were run on each of the scenes. The probability of exceedence plots are shown on Figure 9. From the plots it can be seen that the same improvement in CFAR thresholding is evident. Due to the large number of signatures in the libraries, and the variation between them, the ACE Discriminated detector is shifted further to the right relative to the individual detectors. For this reason, there is a larger deviation between the individual CFARs and the ACE Discriminated Detector CFAR.

11. SUMMARY

In this paper several ACE detection and discrimination pipelines were introduced. The ACE detector and its properties were described, along with how ACE could be used as a discriminator. For the signature dependent and ACE Discrimination thresholding pipelines, it was shown that an accurate CFAR threshold can only be set after discrimination. When one sets a CFAR threshold for each individual detector, the result is a CFAR threshold at a much higher CFAR. This difference between the two CFARs was shown using real hyperspectral data. Using Dataset A, it was shown that ACE discrimination is an effective object classifier.

ACKNOWLEDGMENTS

This work is sponsored by the Air Force Research Laboratory under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

REFERENCES

- [1] Manolakis, D., "Taxonomy of detection algorithms for hyperspectral imaging applications," *Optical Engineering* **44 NO 6**, 066403 (2005).
- [2] Kraut, S. and Scharf, L. L., "Adaptive subspace detectors," *IEEE Transactions on Signal Processing* **49 NO 1**, 1–16 (January 2001).
- [3] Scharf, L. L., [*Statistical Signal Processing*], Addison-Wesley, Reading, MA (1991). ch. 4.
- [4] Shaw, G. and Burke, H., "Spectral imaging for remote sensing," *Lincoln Laboratory Journal* **14 NO 1**, 3–28 (2003).
- [5] Manolakis, D., Marden, D., and Shaw, G., "Hyperspectral image processing for automatic target detection applications," *Lincoln Laboratory Journal* **14 NO 1**, 79–116 (2003).